From: Rex Buddenberg

Reference:

PS Docket No. 11-15

DA 11-175

PUBLIC SAFETY AND HOMELAND SECURITY BUREAU SEEKS COMMENT ON RAPIDLY DEPLOYABLE AERIAL TELECOMMUNICATIONS ARCHITECTURE CAPABLE OF PROVIDING IMMEDIATE COMMUNICATIONS TO DISASTER AREAS

Point of view. I have supervised a number of masters theses in these subject area over the years and am familiar with several communications relay infrastructures in government, including DoD. But the comments are mine personally, not those of any students, employer or clients.

Editorial. I have preserved all your text in the content paragraph of the inquiry unchanged, small font.

The Bureau seeks comment regarding low-altitude

What do you mean by 'low altitude'? 65K feet is low altitude for a satellite but quite high altitude for an air breather.

Environment. In brief, we generally experience weather (which impacts flyability of an airframe, manned or otherwise, tethered or otherwise), from ground level to around 55K feet. Above that – 65K feet is commonly mentioned -- a vehicle is above the jet stream, above cloud cover. It is easy to fly a platform at that altitude – light airs; it's somewhat harder to get it there in the first place. There have been many proposals and experimentals for both winged aircraft and lighter-than-air craft, both in and out of government.

Communications footprint. In layman terms, the horizon distance footprint for a vehicle at 65K feet is about the size of Kuwait.

The spectrum impact here is that we are dealing with radio footprints somewhat larger than commercial terrestrial cellphone cells but smaller than satellite footprinits, so the frequency reuse problem is harder and easier, but not terribly different.

Any protocol expected to operate in this radio-WAN environment must have a contention-free Media Access Controller. The protocol modifications required to reuse either IEEE 802.16 or LTE are confined to the Physical Media Dependent sublayer; the rest of the protocol needs no change[1].

aerial telecommunications architecture .

The term 'architecture' is commonly defined sloppily. And the question you pose invites confusion: does 'architecture' mean the UAV (platform-oriented) infrastructure? Or does 'architecture' mean the communications modularization model (cross-platform)? Or something else? Because building and flying airframes is a quite different discipline than planning communications and internetwork protocols, the consequences of the definitional confusion can be considerable. We have many bad examples of this confusion and failure to define terms.

1 A good share of the requirements work is encapsulated in working group notes of a group organized by USMC and including all the DoD military services and interested vendors.

An 'aerial architecture' can be disposed of fairly easily by defining the interface between the airframe and communications payload. The airframe needs to provide hotel services (weight, cube, electrical power, place to mount antennae, ...) but needs to know little about the functionality of the payload. The remainder of an 'architecture' here then becomes one of provisioning and operating doctrine – what kind of aircraft do you have and where do you fly it?

The 'telecommunications architecture' should consist of a modularization model consisting of routable networks: the aerial platform is hauling one node of a radio-WAN ... the radio-WAN is, in turn, one network segment in an internetwork. This makes the communications protocol selection consistent with the National Broadband Plan. In most cases, the aerial platform would be hauling a base station[2] of a radio-WAN; one or more of the subscriber stations must be attached to a router in the surviving internetwork. The communications node itself is largely indifferent regarding where it happens to be installed.

There are numerous examples, especially in the military, of how not do do this; we needn't describe them, but should note that many of the difficulties root to this confusion of the term 'architecture' and consequent poor modularization.

solutions that are accessible, reliable, resilient, cost-effective and secure, and which are capable of providing public safety and emergency response personnel the capability to communicate during the critical restoration period after a major disaster over a multitude of communications platforms

Providing communications to emergency services personnel is, of course, important. But it is equally important to provide the same communications into disaster areas for citizen use. A solution that only meets one of these requirements is only half a solution. Two infrastructures is an unsustainable situation.

(e.g., HF,

There is little value in elevating an HF (or low-band VHF) antenna. There is value in using HF – if it is covered with a routable network protocol and can fit into the internet. Rapidly deployable, certainly; but rapidly deployable in a UAV is somewhat more problematic and of doubtful value[3].

UHF, VHF, cellular, Internet, and satellite).

The frequency band choice is less important than the frequency-independent protocol choices. The protocol choices should be consistent with the National Broadband Plan's intent of extending the internet. This is because we need to restore internet communications in the disaster area, but leverage the existing internetwork in undamaged parts of the area.

I. You want a protocol that represents a routable network. Because we are reconstituting an internetwork -- router-to-router interconnect. This allows us to focus on restoration of infrastructure while remaining agnostic about applications; we need not have any knowledge of what is on the other side of a router to do the restoration job right.
II. Within that routable network segment, four issues are important to us:
    1. Contention-free MAC. Stability under overload is a critical criteria; bandwidth

2 Base station (BS) is the term per IEEE 802; synonym terms include 'head end' (cable TV) and CPE (telephone).

3 Incorporating HF into an internetwork is an interesting problem. The theory is no different than in higher frequencies and the protocols above Phy should not be materially different, but the difficulty is getting enough contiguous spectrum (200kHz is a practical bare minimum). Proceeding here would require a large scavenging effort similar to the narrowbanding one current in the VHF band.

efficiency and control are secondary positive attributes of a contention-free MAC[4].

2. Point-to-multipoint (aka multicast). The radio-WAN will be the capacity-limited segment of an internet and ability to exploit multicast (e.g. for citizen notifications) is an important offset.
3. Infrastructure protection security measures. The protocol needs to have resistance to denial and theft of service attacks and some jam resistance. (Authenticity and confidentiality are NOT requirements here since they are end-to-end security requirements and can only be solved in applications). Some jam resistance (or interference resistance) measures can be found in protocol construction; some will be found in spectrum management.
4. Simple Network Management Protocol agents should be provisioned in all the equipment. (This is first step to answer your coordination and management questions below and both prerequisite to and independent of operating doctrine). The sufficient reason for SNMP agents is that they, plus a monitoring center, meet the third principle of high availability engineering. But presence of an SNMP management capability offers far more.

We recognize that authorization of any particular aerial
telecommunications architecture for public safety use will depend on various characteristics of that solution,

You have two questions in here, and the failure to define 'architecture' means that we have to sort them out first.

such as the spectrum band in which the equipment operates,

This is one of the questions.

but we are seeking comment here
only to better understand the available technologies as well as the associated technical issues.

The other question, because this inquiry is anticipating aerial vehicles, is how to manage the airspace? You cannot simply fly a UAV or winch up an aerostat into the area – it cohabits with other disaster relief aircraft (e.g. Coast Guard or county sheriff helicopters, Air Force One, ...). The airspace management problem equals the spectrum management one in complexity and the Federal Aviation Administration has procedures that only partially account for UAVs and requirements for rapid reaction[5]. US Army, US Air Force and US Marine Corps all have experience keeping airspace flight-safe. Most of the ones that I've seen are cumbersome, overly complex – cannot be KISS-employed at the time they're needed most (immediately post-disaster).

The Bureau
seeks comment on the means of coordinating and managing such solutions before and during deployment,

The key word in this phrase is 'before'. The aerial component of the internet will be the most expensive to purchase and the most expensive to own and operate[6]. It will also be the least

4 A contention-free MAC solves a critical spectrum management problem quite elegantly. Emergency services clearly need assured access and a contention-free MAC support that. Without requiring separate frequency assignments. Assured access no longer requires segregated frequency bands.

5 This does represent a reason for flying communcations relay UAVs at edge-of-atmosphere altitudes where traffic, especially manned traffic, is sparser.

6 This observation is true in a non-internet communications infrastructure as well. So it does not constitute a valid reason

capacious. An aerial extension of the internet needs to be exercised regularly by the emergency services personnel and their support. There are, just like manned aircraft piloting or other parts of emergency response, some perishable skills that cannot be generated between the onset of a disaster and the need to operate a deployment. These skills fall in the network operations and in the install/configure/troubleshoot categories.

what public safety agencies may be involved in the operations of such systems

There is a myriad of agencies from federal to municipal level. Any agency that operates manned aircraft today my be operating UAVs at some time in the future. Plus some. To an aircraft, the comms-relay function is simply a payload. While this territory may be new to FCC, it is not to FAA.

There are two parts of the communications problem. Each agency may, by choice or default, chose to amalgamate or separate:

- Data relay (or sensor export – essentially the same problem). This, in the military, is called the 'mission package'.
- Avionics control, or flying the aircraft. This represents communications between the flight control on the ground and the UAV.

Both of these problems are communications problems and both could be solved in a unified 'extend the internet' fashion. In practice, they are hardly ever solved in a unified fashion. Most communications satellites today have a payload channel and a satellite control channel. In the case of DoD, whenever a unified approach is even attempted, it is so badly stovepiped to be of little general purpose use.

or those who own
communications operations may be adversely impacted by such systems, specific classes of such systems that might be involved in such operations with their capacity and bandwidth requirements, the costs associated with their design, implementation and deployment and maintenance of such packages.

While there are certainly interference issues, they are not generically different than the ones we are dealing with today. An emitter can interfere with someone else, whether or not it is airborne. The interference problems will be somewhere midway between the easier problem of WiMAX/LTE terrestrial cellphone ones and the harder problem of communications satellite spectrum management.

Since the FCC seems to now be in the protocol-prescription business, the compatibility issues (interference) need to be matched by attention to interoperability ones. No solutions should be considered that are not routable networks (see criteria recited above). This makes the communications system agnostic about the applications that run over it.

Specifically, the Bureau seeks comment on how to best ensure spectrum coordination of these systems with terrestrial and satellite infrastructure in the affected area or adjacent areas, interference mitigation techniques to ensure that terrestrial and satellite communications are not negatively impacted and network management.

The problems are generically similar to ones we already have. My perception is that the spectrum-coordination problems here are little different than the currently-debated White Space issues. And the protocols themselves require require only Phy-layer adaptations of existing IEEE 802.16 and LTE protocols, perhaps with some of the features in IEEE 802.22[7].

for departing from the internet model.

7 The last I examined, nobody was building to the IEEE 802.22 protocol (perhaps pending the White Space decisions in FCC). But the protocol PAR and specification is quite similar to IEEE 802.16.

Further, the Bureau seeks comment on how to ensure that communications that are provided via aerial communications systems for emergency preparedness are secure

The security problem must be divided into two parts. The first part is authenticity and confidentiality of the data (content). Content protection is outside scope of this inquiry as the protection is required end-to-end, not just over the UAV-supported network segment. This first problem must be solved in applications (e.g. NextGeneration911).

The second part of the security problem is infrastructure protection – resistance to denial and theft of service, jam resistance[8]. This part is within scope. Some of these attacks can be made over the internet (i.e. through a router that the radio-WAN interconnects to the rest of the internet with) and some are direct – attack against the network segment itself. None of these attacks are different because a node of the radio-WAN is airborne. So the theft-of-service resistances in the security sublayer of IEEE 802.16 MAC specification are just as applicable to a UAV-based infrastructure as an all-terrestrial one.

The questions you did not inquire about ... but should. The most important part of a disaster-relief internet is not the infrastructure that is deployed into a disaster footprint after a disaster strikes. Rather the most important part is the pre-existing internet infrastructure. The more robust the every-day internet is, the easier, quicker, cheaper and more effective the patching job is.

Recommend. The FCC should balance its inquiry into rapidly deployable aerial infrastructure with:

- An inquiry into how to make the everyday internet infrastructure more disaster-resistant in the first place, making it more easily restored. Building an internetwork that is robust is very straightforward – execute the three principles of high availability engineering. The result is a highly distributed, robust communications system. The centralized electrical power distribution system then becomes the Achilles heel that works against such robustness – in a disaster, lots of perfectly good communications equipment is disabled for lack of local electrical power ... whether or not there is an airborne relay flying over it.
- An inquiry into how to extend/restore the internet into a disaster footprint without resorting to aerial vehicles. Quickly-deployable extend-the-internet equipment that is ground vehicle-mounted or ground vehicle-transported is entirely practical. Such was deployed by colleagues and students in hurricane Katrina. Such is a part of both US Army and US Marine Corps infrastructure today. Many agencies, certainly US Coast Guard, maintain such 'fly-away/drive-away' kits today. This does not eliminate the need for an aerial component but diminishes the size of this inevitably expensive solution.

While both of these prerequisite topics lack the trade show sex appeal of aerial vehicles, they 1) are more economical of the taxpayers' dollars and 2) provide the foundation for an aerial vehicle extension.

Consider, for example, the multitudinous benefits of applying high availability engineering principles (primarily alternate routes, backup (aka off-grid) power, and fault management) to the internet infrastructure already existing in our schools, fire houses and

8 There are military-specific issues dealing with traffic analysis, traffic flow analysis, low probability of intercept that are part of the infrastructure protection discussion but outside of the disaster-relief requirements set.

police stations. One of these benefits is decreasing need to invest in high cost aerial platforms.

These inquiry recommendations are consistent with the National Broadband Plan, especially including Chapter 16.

Thank you.

/s/ Rex Buddenberg